

WL-TM-96-1132

COMPARISON OF SENSOR MANAGEMENT
STRATEGIES FOR DETECTION AND CLASSIFICATION



STAN MUSICK
KEITH KASTELLA

LORAL DEFENSE SYSTEMS-EAGAN

MARCH 1996

FINAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED.

AVIONICS DIRECTORATE
WRIGHT LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT PATTERSON AFB OH 45433-7623

Comparison of Sensor Management Strategies for Detection and Classification[®]

13 March 1996

Stan Musick[†]
Wright Laboratory
Wright-Patterson AFB OH 45433
musicksh@aa.wpafb.af.mil

Keith Kastella^{*}
Loral Defense Systems-Eagar
kkastell@eag.unisysgsg.com
(612) 456-2537

19961125 111

ABSTRACT

Several sensor management schemes based on information theoretic metrics such as discrimination gain have been proposed, motivated by the generality of such schemes and their ability to accommodate mixed types of information such as kinematic and classification data. On the other hand, there are many methods for managing a single sensor to optimize detection. This paper compares the performance against low signal-noise ratio targets of a discrimination gain scheme with three such single sensor detection schemes: the Wald test, an index policy that is optimal under certain circumstances and an 'alert/confirm' scheme modeled on methods used in some radars. For the situation where the index policy is optimal, it outperforms discrimination gain by a slight margin. However, the index policy assumes that there is only one target present. It performs poorly when there are multiple targets while discrimination gain and the Wald test continue to perform well. In addition, we show how discrimination gain can be extended to multisensor / multitarget detection and classification problems that are difficult for these other methods. One issue that arises with the use of discrimination gain as a metric is that it depends on both the current density and an a priori distribution. We examine the dependence of discrimination gain on this prior and find that while the discrimination depends on the prior, the gain is prior-independent.

[®] Paper presented at the 9th National Symposium on Sensor Fusion, 11-13 March 1996, Naval Postgraduate School, Monterey, CA.

[†] Supported by a grant from the Air Force Office of Scientific Research.

^{*} Supported by the Minnesota Center for Industrial Mathematics under Air Force Office of Scientific Research Contract AF/F49620-95-1-0307

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE MAR 1996	3. REPORT TYPE AND DATES COVERED FINAL		
4. TITLE AND SUBTITLE COMPARISON OF SENSOR MANAGEMENT STRATEGIES FOR DETECTION AND CLASSIFICATION		5. FUNDING NUMBERS C PE 61102 PR 2304 TA ES WU 01		
6. AUTHOR(S) STAN MUSICK KEITH KASTELLA				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AVIONICS DIRECTORATE WRIGHT LABORATORY AIR FORCE MATERIEL COMMAND WRIGHT PATTERSON AFB OH 45433-7623		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AVIONICS DIRECTORATE WRIGHT LABORATORY AIR FORCE MATERIEL COMMAND WRIGHT PATTERSON AFB OH 45433-7623		10. SPONSORING/MONITORING AGENCY REPORT NUMBER WL-TM-96-1132		
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Several sensor management schemes based on information theoretic metrics such as discrimination gain have been proposed, motivated by the generality of such schemes and their ability to accommodate mixed types of information such as kinematic and classification data. On the other hand, there are many methods for managing a single sensor to optimize detection. This paper compares the performance against low signal-noise ratio targets of a discrimination gain scheme with three such single sensor detection schemes: the Wald test, an index policy that is optimal under certain circumstances and an 'alert/confirm' scheme modeled on methods used in some radars. For the situation where the index policy is optimal, it outperforms discrimination gain by a slight margin. However, the index policy assumes that there is only one target present. It performs poorly when there are multiple targets while discrimination gain and the Wald test continue to perform well. In addition, we show how discrimination gain can be extended to multisensor / multitarget detection and classification problems that are difficult for these other methods. One issue that arises with the use of discrimination gain as a metric is that it depends on both the current density and an a priori distribution. We examine the dependence of discrimination gain on this prior and find that while the discrimination depends on the prior, the gain is prior-independent.				
14. SUBJECT TERMS			15. NUMBER OF PAGES 20	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR	

GENERAL INSTRUCTIONS FOR COMPLETING SF 298

The Report Documentation Page (RDP) is used in announcing and cataloging reports. It is important that this information be consistent with the rest of the report, particularly the cover and title page. Instructions for filling in each block of the form follow. It is important to *stay within the lines* to meet *optical scanning requirements*.

Block 1. Agency Use Only (Leave blank).

Block 2. Report Date. Full publication date including day, month, and year, if available (e.g. 1 Jan 88). Must cite at least the year.

Block 3. Type of Report and Dates Covered. State whether report is interim, final, etc. If applicable, enter inclusive report dates (e.g. 10 Jun 87 - 30 Jun 88).

Block 4. Title and Subtitle. A title is taken from the part of the report that provides the most meaningful and complete information. When a report is prepared in more than one volume, repeat the primary title, add volume number, and include subtitle for the specific volume. On classified documents enter the title classification in parentheses.

Block 5. Funding Numbers. To include contract and grant numbers; may include program element number(s), project number(s), task number(s), and work unit number(s). Use the following labels:

C - Contract	PR - Project
G - Grant	TA - Task
PE - Program Element	WU - Work Unit Accession No.

Block 6. Author(s). Name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. If editor or compiler, this should follow the name(s).

Block 7. Performing Organization Name(s) and Address(es). Self-explanatory.

Block 8. Performing Organization Report Number. Enter the unique alphanumeric report number(s) assigned by the organization performing the report.

Block 9. Sponsoring/Monitoring Agency Name(s) and Address(es). Self-explanatory.

Block 10. Sponsoring/Monitoring Agency Report Number. (If known)

Block 11. Supplementary Notes. Enter information not included elsewhere such as: Prepared in cooperation with...; Trans. of...; To be published in.... When a report is revised, include a statement whether the new report supersedes or supplements the older report.

Block 12a. Distribution/Availability Statement. Denotes public availability or limitations. Cite any availability to the public. Enter additional limitations or special markings in all capitals (e.g. NOFORN, REL, ITAR).

DOD - See DoDD 5230.24, "Distribution Statements on Technical Documents."

DOE - See authorities.

NASA - See Handbook NHB 2200.2.

NTIS - Leave blank.

Block 12b. Distribution Code.

DOD - Leave blank.

DOE - Enter DOE distribution categories from the Standard Distribution for Unclassified Scientific and Technical Reports.

NASA - Leave blank.

NTIS - Leave blank.

Block 13. Abstract. Include a brief (*Maximum 200 words*) factual summary of the most significant information contained in the report.

Block 14. Subject Terms. Keywords or phrases identifying major subjects in the report.

Block 15. Number of Pages. Enter the total number of pages.

Block 16. Price Code. Enter appropriate price code (*NTIS only*).

Blocks 17. - 19. Security Classifications. Self-explanatory. Enter U.S. Security Classification in accordance with U.S. Security Regulations (i.e., UNCLASSIFIED). If form contains classified information, stamp classification on the top and bottom of the page.

Block 20. Limitation of Abstract. This block must be completed to assign a limitation to the abstract. Enter either UL (unlimited) or SAR (same as report). An entry in this block is necessary if the abstract is to be limited. If blank, the abstract is assumed to be unlimited.

1. Introduction

The problem of sensor management is to determine how to select sensors, sensor modes and sensor search patterns to maximize the effectiveness of individual sensors and collections of sensors which may be located on different platforms against a set of mission requirements [Musick, Popoli]. For example, a typical air combat application entails the management of onboard radars, electro-optical sensors and passive radar sensing devices. Each of these sensors has many modes to choose from. Radars use different wave-forms and scan patterns for search, raid assessment, target identification and fire control while imaging sensors may have several spectral bands and magnifications to choose from. Systems collect several types of information such as position and classification data. For example, consider the coordination of a radar with an imaging sensor such as a Forward Looking Infrared (FLIR) camera. The FLIR provides azimuth and elevation information on the target which can be used for kinematic tracking, in addition to the FLIR image which can be used to classify the target. This diversity of information must be accounted for in a sensor manager.

Many types of approaches have been developed that address parts of the overall problem of sensor management. Some of the work in this area has focused on problems where the data are primarily kinematic with little emphasis on the detection/classification element of the problem. A Kalman filter can be used to predict how target state estimates evolve depending on variables such as estimated target dynamics, range, and signal/noise ratio [Blackman86, Nash, Watson]. The update rate required to maintain the target position variance at some desired value can be determined by the trace of the position components of the covariance estimate. As target dynamics increase, the required update rate also increases. A threshold in the dynamics can be used to switch the radar to different modes such as a higher energy waveform when it is appropriate. This adaptive approach yields substantial performance improvement over nonadaptive sensor methods by allowing targets to be tracked with reduced time-on-target leading to higher target capacity for a given sensor [Watson]. However, it does not easily generalize to situations where detection, tracking and classification are being performed by the same set of sensors and must simultaneously be optimized.

Another type of sensor management that is used on some existing radar systems is the so-called 'alert/confirm' method [Blackman92]. With this, a portion of the surveillance volume is scanned. If an initial tentative detection is obtained at some point (an 'alert'), then the radar dwells on that point in order to confirm the target location.

There is also a mature literature from the field of operations research on the related problem of optimal search [e.g. Stone]. These applications treating the problem of how to direct a sensor to find an object (which can be fixed or moving [Mangel]) constitute a subset of the general problem of sensor management. Analytical results have been obtained for some cases such as certain types of constraints on the search patterns, random fluctuations in the sensor sweep width, and the inclusion of some types of false alarms [Stone].

Recently, [Castanon] used stochastic dynamic programming to analyze optimal search for detecting a single static target located in one of C detection cells. Remarkably, under a symmetry condition on the measurement probability density, the probability to detect the target is maximized by an index rule. This index rule is to select the cell that is most likely to contain the target for each sensor dwell. What is especially interesting about this result is that greedy (or myopic) optimization, i.e. selecting each sensor dwell to maximize the immediate gain, yields the global optimum.

Optimality proofs along the lines of [Castanon] are quite useful for guiding the selection of sensor management strategies, even though when they apply only under somewhat restricted circumstances. More general strategies are often required to handle realistic assumptions such as multiple targets and the need to simultaneously detect and classify targets. Such general strategies may be suboptimal if they are not strict generalizations of an optimal policy such as this index policy. However, a good general strategy should perform nearly as well as the optimal strategy in the restricted case.

Castanon also examines some problems that violate the symmetry requirement for which myopic optimization does not yield the global optimum. For a collection of small randomly generated problems whose optimum can be found by exhaustive enumeration, [Castanon] shows that the greedy index policy performs nearly as well as the globally optimal solution and somewhat better than an entropy maximization approach.

In order to simultaneously optimize conflicting objectives such as detection, tracking and classification, several authors have proposed the use of measures derived from information theory [e.g. Hintz]. This can be used to coordinate the collection of many different types data and shifts the emphasis from optimizing parameter estimates for individual targets to optimizing the probability density estimates constructed by data fusion systems. By measuring the effect of sensor observations on the probability densities that are used to represent knowledge of a scene, information theory can weigh the tradeoff between different types of information such as positions, velocities and attributes in a single integrated performance measure. This can accommodate tradeoffs between sensors that measure different aspects of a scene, instead of limiting the focus to optimizing target position estimates.

One information theoretic metric that has been applied to several types of sensor management problems is discrimination gain. This has been used for multisensor / multitarget assignment problems [Schmaedeke93], minimizing error correlations between close targets [Kastella94], single sensor detection / classification problems [Kastella95]. This entails predicting the expected discrimination gain for each sensor dwell which is similar to covariance prediction in a Kalman filter. For a simple detection problem discrimination gain results in a performance improvement on the order of 6 dB compared with direct search [Kastella97]. These types of performance increases are significant if similar results hold for realistic systems.

This paper is organized as follows. The next section defines several metrics for sensor management. We emphasize discrimination gain (DG) and show how it can be used for multisensor / multitarget detection and classification problems. One issue with using discrimination as a basis for sensor management is that it depends on an a priori distribution. There are several alternatives for this prior, so it is natural to ask how this prior affects the discrimination gain here. Also, an alternative metric based on optimization of the probability to correctly determine the state of each cell in the surveillance volume is also developed. Section 2 concludes with definitions of three other sensor management metrics for detection to be used for performance comparison with DG. Section 3 examines comparative performance for single and multitarget detection with a single sensor. DG, alert/confirm and the index rule of [Castanon] are evaluated using a monte carlo simulation with discrete measurements. Detection performance for a sequential probability ratio test is estimated using an analytic expression for expected time to decision [Wald, Blackman86]. In Section 4 the behavior of DG for multitarget / multisensor detection / classification is compared to greedy optimization of the total probability metric. The time-dependence of the error rates and sensor use are obtained using monte carlo simulation. Surprisingly, direct optimization of the total probability metric saturates quickly leading to very poor performance. Section 5 discusses the results and their implications.

2.0 Sensor Management Strategies

Five detection/classification techniques are developed. Their short names are direct search, alert/confirm, index rule, SPRT (sequential probability ratio test), and DG (discrimination gain). All five techniques are normative, all are sequential, and all but SPRT use Bayesian methods to process the measurement data recursively. By normative we mean based on principled probabilistic models and employing some form of mathematical optimization scheme. What distinguishes these techniques from each other is mainly the policies they adopt for sequencing between cells. Although these techniques are quite different in terms of their search policies, each is a reasonable choice worthy of consideration here. The older techniques like SPRT have proven themselves in a variety of diverse applications, while the newer techniques have shown good promise in research studies [Castanon, Kastella97, Jenson].

Consider the following problem. There are an unknown number of targets confined to a surveillance volume that consists of C discrete cells indexed $c = 1, \dots, C$. Each cell contains at most one target: it can be empty or it can contain a target that from one of T target classes. The state of each cell is then indexed by $t = 0, \dots, T$ where $t = 0$ denotes an empty cell. The problem of detection and classification is then to determine the state $t = 0, \dots, T$ of each cell.

There are S sensors labeled $s = 1, \dots, S$. The cells are sampled at discrete times k . At each time k a single cell is sampled. Only one sensor can be used at each sample time. The probability densities $p(z|t, s)$ to obtain measurement outcome z given that sensor s is used and that the cell contains target type t are known, time-independent and independent of c . The individual measurements z can consist of continuous, discrete, vector- or set-valued random variables. The expression $\int dz f(z)$ shall denote the integral over the entire measurement domain. Let N^k be the number of measurements in a cell c at time k . The entire set of measurements in cell c is $Z = \{z_1, \dots, z_{N^k}\}$. In general, Z contains measurements produced by a number of different sensors.

Our objective is to compute the conditional density $p(t|Z, c)$ for cell c to contain a target of type t , given the observation set Z . Then the minimum-probability-of-error classifier is to classify each cell according to [Blahut]

$$\hat{t}_c = \arg \max_t p(t|Z, c). \quad (1)$$

The error probability for that cell is then $p_e(c) = 1 - p(\hat{t}|Z, c)$.

Suppose that sensor s is used to perform a measurement in the cell, producing outcome z . The new set of observations is $Z' = \{z\} \cup Z$ and the new target probability is obtained using Bayes rule,

$$p(t|Z', c) = \frac{p(z|t, s)p(t|Z, c)}{\sum_{t'} p(z|t', s)p(t'|Z, c)}, \quad (2)$$

where we have assumed that the cell-state hypotheses are independent. The a priori density is determined by the relative frequencies of occurrence for each target type. The outcome of a measurement can be predicted, conditioned on the current measurement set as

$$p(z|Z, c, s) = \sum_{t=0}^T p(z|t, s) p(t|Z, c). \quad (3)$$

2.1 Discrimination Gain Metric

Let $p(t)$ and $q(t)$ be two distributions. Then in general, the discrimination between them is [Blahut, Kapur]

$$L(p; q) = \sum_{t=0}^T p(t) \log(p(t) / q(t)). \quad (4)$$

For the application at hand, let $p_0(t|c)$ be some a priori probability for target type t to be in cell c . Expanding Eq. (4), the discrimination in cell c is

$$L(Z, c) = \sum_t p(t|Z, c) \log(p(t|Z, c)) - \sum_t p(t|Z, c) \log(p_0(t|c)). \quad (5)$$

The expected discrimination when one additional measurement is made in cell c can be computed using the density $p(z|Z, c, s)$. Using Eqs. (2-5), the expected discrimination is

$$\begin{aligned} E_{Z, c, s}[L(c, Z')] &= \int dz p(z|Z, c, s) L(p(t|Z', c); p_0(t, c)) \\ &= \int dz \sum_t p(z|t, s) p(t|Z, c) \log(p(t|Z', c)) \\ &\quad - \sum_t p(t|Z, c) \log(p_0(t|c)) \end{aligned} \quad (6)$$

The interesting feature to observe here is that the prior-dependent terms of $L(Z, c)$ and $E_{Z, c, s}[L(c, Z')]$ are identical, so they cancel in the evaluation of $\Delta L(Z, c, s) = E_{Z, c, s}[L(c, Z')] - L(c, Z)$,

$$\begin{aligned} \Delta L(Z, c, s) &= \int dz \sum_t p(z|t, s) p(t|Z, c) \log(p(t|Z', c)) \\ &\quad - \sum_t p(t|Z, c) \log(p(t|Z, c)) \end{aligned} \quad (7)$$

As a result, the discrimination gain is independent of the prior. ($\Delta L(Z, c, s)$ is also the same as the decrease in entropy.)

2.2 Direct Optimization of Total Probability to Correctly Classify

As an alternative to DG, the probability that all of the cell-states are estimated correctly can be maximized directly for each cell sample. With the maximum likelihood classifier Eq. (1) the probability to correctly classify cell c is

$$p_{corr}(Z, c) = 1 - p_e(c) = \max_t p(t|Z, c) \quad (8)$$

Using Eq. (3), the expected increase in $p_{corr}(c)$ for the individual cell when it is sampled with sensor s is $\Delta p_{corr}(Z, c, s) = E_{Z, c, s}[p_{corr}(Z', c)] - p_{corr}(Z, c)$ where $Z' = \{z\} \cup Z$ and

$$\begin{aligned} E_{Z, c, s}[p_{corr}(Z', c)] &= \int dz p(z|Z, c, s) \max_t [p(t|Z', c)] \\ &= \int dz \max_t [p(z|t, s) p(t|Z, c)]. \end{aligned} \quad (9)$$

The probability to correctly classify all of the cells in the surveillance volume is

$$p_{corr}^{tot} = \prod_{c=1}^C p_{corr}(Z, c). \quad (10)$$

When cell c is observed with sensor s , then the incremental increase in p_{corr}^{tot} is

$$\Delta p_{corr}^{tot}(s, c) = \left(\prod_{c'=1}^C p_{corr}(Z, c') \right) \frac{\Delta p_{corr}(Z, c, s)}{p_{corr}(Z, c)}. \quad (11)$$

In Eq. (11) the prefactor $\prod_{c'=1}^C p_{corr}(c')$ does not depend on which cell is observed so that maximizing the gain

$$g(Z, s, c) \equiv \frac{\Delta p_{corr}(Z, c, s)}{p_{corr}(c, s)} \quad (12)$$

is equivalent to maximizing p_{corr}^{tot} .

2.3 Sensor Management for Detection

Let us postulate a situation where there are just two possible outcomes from a sensor, either detection or no detection. Single-observation detection and false-alarm probabilities are denoted

$P_D = p(z = 1|t = 1)$ and $P_F = p(z = 1|t = 0)$, where cell type 0 represents no target and cell type 1 represents target. A slightly more general condition is when there are $T > 1$ target types and the full set of types is denoted $\{0, 1, 2, \dots, T\}$. Assuming the sensor can distinguish between $T+1$ types, the measurement density matrix ($\Pr(\text{sense type } \tilde{t} | \text{actual type } t)$) is square and of order $T+1$. In the performance comparisons in Section 3, the focus will be on $T=1$ and on dim targets where P_D is 0.69 and P_F is 0.31.

To compare results between the various detection techniques, let c_t denote the collection of cells that contain a target of type t . Further denote the actual target type in a particular cell c as t_c . We define two error probabilities as follows:

$$\begin{aligned} p_e(k) &= \Pr\{\arg \max_c p(t=1|Z_k, c) \in c_0\} \\ p_{et}(k) &= \Pr\{\max_{\tau \neq t_c} p(\tau|Z_k, c_t) > p(t_c|Z_k, c)\} \end{aligned} \quad (13)$$

p_e is the frequency for the largest probability in the cell volume occurring in a non-target cell. p_{et} is the frequency in type t cells for the largest probability to be for a type t^* other than t . If $T=1$, the only types present are non-targets and targets. Then $p_{e0} = p(\hat{t}=1|t=0)$ is the frequency in non-target cells for the largest probability to be a target while $p_{e1} = p(\hat{t}=0|t=1)$ is the frequency in target cells for the largest probability to be a non-target. p_{e0} and p_{e1} are independent of one another and are not related by formula to p_e . Here p_e is a global error metric in the sense that it looks at the entire volume to produce a single decision while p_{et} is local since it focuses on a representative cell of type t . It is straightforward to construct expressions for these probabilities in terms of the errors arising in simulation.

We now briefly describe several alternative sensor management techniques that will be compared to DG in single sensor detection problems.

Direct search (DS): The search procedure in the direct search technique is to advance through the cells in the volume in the same order for every frame, taking one measurement in each cell, and repeating the frame as many times as are necessary, say to reach a decision. This is the simplest of the five techniques because no decisions are required in order to gather measured data. Even though the individual measurements are processed optimally, we should not expect direct search to be a high performer but it is included here to provide a baseline for comparisons.

Alert/confirm (A/C): This technique proceeds exactly as direct search does until some cell yields a detection called an alert [Blackman92]. An alert triggers a confirm cycle in which additional measurements are immediately taken in the alert cell. The dwell time for the additional measurements is often longer than that for alert so that a higher signal/noise ratio (SNR) can be achieved. For a radar using coherent integration, the SNR is proportional to the dwell time. A variety of mechanizations are available here but short-dwell/long-dwell is reasonable in many applications and that is what we will assume.

Index rule (IR): The index rule's search procedure is to look in only the most likely cell, that is the cell with the highest probability of containing a target. [Castanon] shows that this greedy procedure is optimal for all time if two key assumptions are true. First, the measurements z are scalars with the symmetry condition $p(z|t=0) = 1 - p(b-z|t=1)$ for some constant b . Symmetry of this kind often holds for detection problems where observations are modeled as a known signal plus noise, the noise itself modeled as a symmetric Gaussian density (for binary measurements this symmetry becomes $P_D = 1 - P_F$). The second assumption is that the search volume contains just one target. Our problem formulation usually satisfies the first assumption but often violates the second. Results will be shown later for both cases, the second assumption satisfied and then violated.

Discrimination gain (DG): This technique is based on a recursive expression for calculating the expected discrimination gain $\Delta L(Z, c, s)$ where an observation with a given measurement density is to be taken in cell c where the current probability vector is $p(t|Z, c)$. The search procedure is to measure the cell c^* with the largest $\Delta L(\cdot)$. After the measurement is made and the Bayesian update to $p(t|Z, c^*)$ is computed, the update is used to re-evaluate $\Delta L(c^*)$. The cells are maintained in a priority queue determined by their expected gains. The computational overhead of this queue is $O(\log C)$ where C is the total number of cells.

Sequential Probability Ratio Test (SPRT): SPRT [Wald] is a well-known technique for binary hypothesis testing. It can be applied to the search problem by evaluating the volume one cell at a time, deciding the correct hypothesis in each cell before moving to the next. After a cell measurement is made and assimilated by the SPRT algorithm, the decision procedure is as follows: 1) accept the hypothesis, 2) reject the hypothesis, 3) defer the decision and take another measurement. Notice that SPRT is the only one of our five techniques that provides explicit criteria for making a final decision about the content of a cell. An algorithm for SPRT is described in [Blackman86]. It takes as input P_D , P_F , and two allowable error probabilities, α and β :

α = probability of rejecting H_0 when H_0 is true
 β = probability of accepting H_0 when H_0 is false

Blackman86 shows the following equation for the expected number of measurements required to reach a decision:

$$E[K_{decide} | \theta] = \frac{P(\theta) \ln((1 - \beta) / \alpha) + [1 - P(\theta)] \ln(\beta / (1 - \alpha))}{d(\theta)} \quad (14)$$

where

$$\theta = \begin{cases} \text{hypothesis } H_0 \text{ (false target) correct} \\ \text{hypothesis } H_1 \text{ (true target) correct} \end{cases}$$

$$d(\theta) = \begin{cases} P_F a_1 - a_2, & H_0 \text{ correct} \\ P_D a_1 - a_2, & H_1 \text{ correct} \end{cases}$$

$$a_1 = \ln \left(\frac{P_D / (1 - P_D)}{P_F / (1 - P_F)} \right), \quad a_2 = \ln \left(\frac{(1 - P_F)}{(1 - P_D)} \right)$$

Now consider an entire volume of cells in which the fraction of targets is $p(H_1)$. Then the expected number of measurements required to reach a decision in a randomly selected cell is

$$E[K_{decide}] = (1 - p(H_1)) E[K_{decide} | H_0] + p(H_1) E[K_{decide} | H_1] \quad (15)$$

The expected number of measurements to reach a decision will become a handy metric for comparing SPRT to the other techniques.

3. Detection Performance Comparison

In this section we examine results for a test case with each detection technique. The evaluation methodology is Monte Carlo simulation, except for SPRT where approximate quantities for comparison can be derived analytically. The primary evaluation criteria are the various error probabilities, a fair basis for comparison that is equally applicable to all techniques. The definition of error probability was given in Eq. (13).

The standard test problem (STP) assumes that a volume of $C=100$ cells must be searched for stationary targets. Each cell contains at most one target, and the a priori distribution for a target to be in a given cell is uniform and known. Initially we assume just one target type is present. Thus each cell is either empty ($t_c=0$) or contains a type 1 target ($t_c=1$). The hypothesis under test in each cell is H_0 : there is no target present in this cell. We further assume a single agile sensor that makes measurements in a fixed amount of time, regardless of cell location. Finally, we assume that the sensor is operating at peak power against exclusively dim targets that produce returns with signal-to-noise ratios of -3dB. At an SNR of -3dB, $P_D = 0.69$ and $P_F = 0.31$, values that will serve for comparative evaluations in this section unless stated otherwise.

We must also consider how the confirm cycle of alert/confirm differs from alert. If confirm uses more resources than alert, e.g. uses more time or power, the simulation results must somehow reflect that fact so all techniques are treated evenly. We will assume to start off that alert and confirm use equal amounts of the sensor resource. Consequently, they have the same P_D and P_F . Later we will allow the confirm cycle to be used more effectively by assuming that P_D builds to 0.9 and P_F drops to 0.1. For a Gaussian target $P_D = 0.9$ corresponds to an SNR of +5.2dB. With coherent processing of pulse returns, producing a confirm measurement +8.2dB higher than alert would require the dwell time to be roughly $6.6 (10^{0.82})$ times longer than that of alert. The extra time required for the confirm cycle will be reflected in presenting graphical results, as discussed later.

A tracking radar mated with an electronically-scanned antenna and operating in its velocity-search and range-while-search modes is representative of the kind of sensor we have in mind here. However, the assumptions of the STP constitute an idealized problem that has been developed to study research issues. We do not claim that a radar designed to the results shown here would constitute a desirable system so caution must be exercised in interpreting these results.

3.1 Simulation Results

This section presents comparisons of the five search techniques based on Monte Carlo simulation results. Comparisons will be drawn using the appropriate performance metric from the error probabilities defined in Eq. (13). Note that each Monte Carlo study consisted of 1000 runs with 1000 measurements (100 cells x 10 measurements/cell on average) taken in each run for a total of 1 million measurements processed per study.

Figure 1 shows comparative results for all four techniques evaluated using the global metric p_e for the STP. The x axis of Figure 1 (and all subsequent figures) is discrete time $k' = k/C$ where k is the total number of measurements processed and C the (fixed) cell count. Note that output sampling occurs only at integral values of k' from 0 to 10, i.e. after a set of C observations is collected. This sampling rate is not a limitation of any aspect of this formulation but merely a convenience for producing output snapshots.

Since all conditions for optimality of the index rule are satisfied in the STP, we should expect the index rule to perform best and it does. The performance of DG is a close second to the index rule.

Figure 1 shows that performance for direct search is worst of the four techniques, as expected; adaptive techniques should outperform the non-adaptive one and they all do. The fifth curve in Figure 1 corresponds to direct search against a stronger target of +3dB SNR. Since this curve is closely aligned with the curves for the index rule and DG, we conclude that these smart techniques have achieved approximately a 6dB gain over direct search.




Figure 1. Comparison of three search methods against optimal index rule.

The last of the four simulated techniques shown in Figure 1 is alert/confirm. It performs slightly better than direct search but not nearly as well as the index rule or discrimination gain. The primary reason for its poor performance may be that it does not dwell on any object (target or no) long enough to resolve it. For this phase of study, power on target in confirm equals that in alert and evidently the slight extra exploration of detected cells is insufficient to actually produce a great improvement. One indicator of how improved performance comes about is the percentage of resources devoted to the target cell. Whereas the index rule expended 23.7% of its time there and DG 18.7%, A/C spent only 1.15% of its time searching the target cell. Although A/C has presumably attempted to adaptively apportion its sampling, its sensor usage is little different from that for direct search where usage is exactly 1%. Another indicator of the similarity of A/C and direct search is the fact that their average α error rates are quite close as are their average β error rates. We now examine how A/C performs if the confirm cycle dwell time is increased.

Figure 2 presents results for A/C when $P_D = 0.9$ and $P_F = 0.1$ on the confirm cycle. The metric in Figure 2 is again the global metric p_e . Figure 2 shows that performance of the optimal index rule and raw A/C are quite close, IR being slightly better at first and A/C slightly better at the end. A/C has the potential to be better than optimal IR because it expends more resources than optimal does. Statistics available at run conclusion show that A/C was in alert status 76.2% of the time (1 unit) and in confirm the remainder (6.6 units). Therefore the actual time required to achieve the performance displayed by A/C in Figure 2 is $.762(1) + (1-.762)(6.6) = 2.33$ times longer than IR took. The curve labeled "alert/confirm effective" in Figure 2 is the raw A/C curve "stretched" along the x axis by the 2.33 factor.

Effective A/C performance is only slightly better than it was when the dwell time was short in confirm, and it still falls considerably short of optimal or discrimination gain results. Sensor use against the target cell increased to 1.29%, still far short of DG and IR values.




Figure 2. Comparison of enhanced alert/confirm to optimal index rule.

Figure 3 compares DG and IR when IR's single-target assumption is violated. Here five targets are scattered randomly through the surveillance volume. Once IR finds a target cell, it remains focused on that cell, refusing to search for other targets. The performance metric for the figure is the local metric p_{et} which increases when targets are not found. The behavior of the IR is optimal in many problems (e.g. when looking for a single-point failure in a broken machine) but not optimal in search when multiple targets are likely to be present. DG is the superior choice for this set of circumstances, as shown in the figure.




Figure 3. Performance of optimal index rule when target count condition

Figure 4 presents error metrics for the DG method. Figure 4a is the metric corresponding to α ($=p_{e0}$) and Figure 4b that for β ($=p_{e1}$). For a particular time k , one may enter these two curves and pick off the (α, β) pair for that k . When an (α, β) pair is coupled with P_D and P_F , all information is available for computing the expected number of samples that SPRT would take to reach a decision, Eq. (15). That was done for the 10 (α, β) pairs in Figure 4. The result is the two curves of Figure 5, one for SPRT and one for DG. The SPRT curve is what Eq. (15) predicts, the DG curve (a straight line) is what the simulation produced. Figure 5 shows that SPRT does slightly better than DG through the first two measurements but then DG dominates from there on out. DG does not enjoy a large or growing advantage over SPRT but its edge does exist. Any of the four techniques could be compared to SPRT in this manner and the rank ordering of results would not change from what we have already seen.




Figure 4a. Experienced error rate in non-target cells, disc gain search.




Figure 4b. Experienced error rate in target cells, disc gain




Figure 5. Comparison of expected time to decision, SPRT and DG.

4. Detection / Classification Results

Typical performance of DG and $\Delta p_{corr}(Z, c, s)$ for detection and classification in a two sensor, multitarget application is shown in Figures 6-10. It should be emphasized that these results are obtained using greedy optimization: each sensor sample is selected to maximize the immediate gain. Figures 6-8 show how the average of the confusion matrix varies as a function of the average number of samples per cell. The confusion matrix $P(\hat{t}|t)$ is the probability to declare that a cell state is \hat{t} given its true state t . In all cases there are 18 targets of class 1 and 2 targets of class 2 (recall the $t = 0$ denotes an empty cell). Figures 6,8 show results with 100 cells in the surveillance volume while Figure 7 has 1000 cells. The curves give the average performance obtained over 500 trials.

The sensors are characterized by the conditional probabilities $p(z|t, s)$. The output z of the sensor takes one of three discrete values, $z = 0, 1$ or 3 with probability determined by t and s . The sensors are selected so that $s = 1$ represents a 'detection sensor' while $s = 2$ is a 'classification sensor'. This can be achieved by defining the sensor measurement density matrices as

Detection Sensor ($s = 1$):

$p(z t,1)$	$t = 0$	1	2
$z = 0$.8	.2	.2
1	.1	.4	.4
2	.1	.4	.4

Classification Sensor ($s = 2$):

$p(z t,2)$	$t = 0$	1	2
$z = 0$.5	.5	.5
1	.25	.4	.1
2	.25	.1	.4

The detection sensor has no ability to distinguish between targets of class 1 and 2 but has detection performance corresponding to about 5 dB target SNR. On the other hand, the classification sensor provides very little information when no target is present ($t = 0$), but when a target is present, it is unlikely to confuse Class 1 with Class 2. Sensor use by type is shown in Figures 9-10.

The missed detection performance is shown by the curves for $P(0|1)$ and $P(0|2)$ in Figures 6-8. For DG these are similar to those obtained for the pure detection problems shown in Figure 1. The classification performance is indicated by the other curves. As expected, the missed target detections are monotone decreasing. Notice that the curve for Class 2 significantly lags the curve for Class 1. This is because there are many more Class 1 targets. Therefore, as non-empty cells are detected, they are initially classified as Class 1 targets based on the a priori information. It is only after several dwells have been performed with the classification sensor that Class 2 targets can be separated from the Class 1 targets. This is apparent from the behavior of the $P(1|2)$ error curve which actually increases initially, as a result of the missclassifications caused by the prior.

Interestingly, this confusion between Class 1 and Class 2 is reduced as the number of cells increases, as shown in the 1000 cell example of Figure 7. The detection performance is quite similar for the 100 and 1000 cell cases. However, once the non-empty cells are detected, a fixed number of Sensor 2 dwells suffices to differentiate between Class 1 and Class 2. Therefore, performance is better as a function of average samples per cell.

Further insight into the behavior of the DG algorithm is provided by Figure 10, showing the time dependence of the sensor use, again plotted as a function of average samples per cell. For the first several dwells, Sensor 1 is used exclusively. It is only after the existence of a target has been established that any benefit to be derived from Sensor 2. This is a desirable behavior for the algorithm since it means that dwells with Sensor 2 are not wasted on empty cells. In fact, the probability to sample an empty cell with Sensor 2 is 10^{-2} for the 100 cell problem and 10^{-3} for the 1000 cell problem.

Surprisingly, greedy optimization of the probability correct using g (Eq. (12)) leads to very poor performance as shown in Figure 8 for the 100 cell, 2-sensor, 2-target class case treated above using DG. The source of this difficulty can be seen by plotting the expected gains for the two sensors as a function of probability vector ($p(t = 0), p(t = 1), p(t = 2)$) =

$(1 - p_t, p_t r, p_t (1 - r))$. With this parameterization the variable p_t is the probability that there is a target in the cell and r is the relative probability that the target is of Class 1, given that there




Figure 6. Time-dendence of confusion matrix $P(\hat{t}|t)$ using DG for two sensor, two target class problem with 100 cell surveillance volume.




Figure 7. Time-dendence of confusion matrix $P(\hat{t}|t)$ using DG for two sensor, two target class problem with 1000 cell surveillance volume.




Figure 8. Time-dendence of confusion matrix $P(\hat{t}|t)$ using probability correct metric for two sensor, two target class problem with 100 cell surveillance volume.




Figure 9. Time-dendence of sensor use for 100 cell problem using discrimination gain of Figure 6. The detection sensor S_1 is used exclusively for the first part of the search. Later, the classification sensor S_2 is used.




Figure 10. Time-dendence of sensor use for 100 cell problem of Figure 8. The expected gain for both sensors saturates at 3 samples/cell (average). Beyond 3 samples/cell, the algorithm selects randomly between sensors.




Figure 11. Expected gain as a functions of probability that a target is present p_t and relative probability r that target is is type 1 or type 2, using discrimination gain and probablity gain metrics.

is a target in the cell. Surface plots for both sensors of the discrimination gain ΔL (Eq. (7)) and probability gain g (Eq. (12)) as functions of p_t and r are shown in Figure 11. Notice that the surface for the detection sensor S_1 is peaked at $p_t = .5$ and is independent of r . On the other hand, the gain using the classification sensor S_2 is maximal at $p_t = 1$ and $r = .5$, corresponding to the situation where one knows for sure that the cell contains a target but one is completely uncertain as to its type.

While the behavior of the probability gain g is qualitatively similar to the discrimination gain, notice that for both sensors, g is 0 over the regions near $p_t = 0$ and $p_t = 1, r = 1$ or 0. For these regions, the probability gain is the same for both sensors, so the metric provides no guidance about which one to use. Once enough data have been collected to place all of the cell probabilities into one of these regions, the algorithm can do no better than to switch between S_1 and S_2 randomly. This behavior is exhibited in the sensor use plot for g shown in Figure 10. The probability to use either sensor quickly goes to .5. This contrasts with the situation for discrimination gain, where the gains are the same only along lines. As a result, there is almost never any ambiguity as to which sensor provides the higher gain.

5. Discussion

This paper has compared several techniques for managing sensors to detect and classify dim targets. Because of its flexibility, robustness, and near optimality, discrimination gain (DG) appears to be the best of the techniques studied. For the restricted problem of detecting a single target, DG and sequential probability ratio test (SPRT) perform similarly while the index rule enjoys a moderate performance advantage, as long as there is only a single target. The advantage of DG is that its generalization to multiple targets, multiple target classes and multiple sensors is straightforward.

When compared to SPRT, DG finds targets slightly more quickly for the same error probability. However, in SPRT, the distribution of time to decision (K_{decide}) is skew [Wetherill, Chs 4,6] and relatively flat. This can cause SPRT to have long test sequences. In applications, this is partially cured by using a truncated sequential probability ratio test. DG may have smaller variance in the time to decision because it incorporates the prior and uses Bayesian updating. Furthermore, SPRT methods are cumbersome for implementing M-ary hypothesis tests ($T > 1$ target types) whereas DG extends naturally to this case.

In [Castanon], a probability based metric treats a detection problem similar to [Kastella97]. [Castanon] provides an optimal detection strategy (presented as an index rule) that holds as long as there is only one target and subject to certain constraints on the sensor properties. Similar probability-based metrics can be obtained in the multitarget/multisensor case, but have not been previously studied. It may be possible to obtain optimality criteria for these metrics using a similar approach. For the case of unknown target number, the optimal strategy resembles a maximum information gain approach. An interesting question is whether the two methods are equivalent in the large sample limit. In relation to the optimal index rule of [Castanon], our

numeric results show that DG is a close second in performance when there is only one target in the search volume. When there are many targets, DG finds them all and is easily the better performer. This leads us to conclude that DG is more robust.

In view of the similar performance of DG, SPRT, and the index rule, it might be expected that most reasonable sensor management strategies will work well. Therefore, it is somewhat surprising that the probability gain metric performs so poorly. This may be due to the fact that greedy optimization was used, suggesting that for this case, it yields solutions that are far from the global optimum. This contrasts with the result of [Castanon], where even when it is not optimal, the greedy index policy yields a very good solution in the single target detection problem.

In comparison with alert/confirm (A/C), DG performs significantly better by every metric we studied, even though we assume coherent integration for the confirm dwell, while DG is essentially an incoherent technique. We believe the primary reason for A/C's poor performance is its simplistic decision policy; i.e. every detection causes an additional measurement to be taken in the cell that produced the detection. In effect, this policy wastes measurements because often the cell contains no target and the "decision" can be made more quickly. The tradeoff here is between the additional complexity of DG for deciding where to search and the lost performance of A/C for not deciding. We believe this supports the case for considering DG in detection/classification applications in radars and other sensor systems where A/C or A/C-type techniques have traditionally been used. However, we are not stating that DG is strongly preferred to A/C. First, both techniques are robust and flexible. Second, we have not studied this problem exhaustively; e.g. only one "point" in the operating curve has been examined. Finally, implementation considerations (such as compute load) may strongly favor one or the other and we have not examined these at all.

Acknowledgment Thanks are due to T. Jenison for his assistance with the numerical experiments in classification and detection performance.

REFERENCES

[Blackman86] Samuel S. Blackman, *Multiple-Target Tracking with Radar Application*, Artech House, Norwood, MA, 1986.

[Blackman92] S. S. Blackman, "Multitarget Tracking with an Agile Beam Radar", in *Multitarget-Multisensor Tracking: Advanced Applications, Vol. II*, Y. Bar-Shalom, Ed., Artech House, 1992.

[Blahut] R. E. Blahut, *Principles and Practice of Information Theory*, Addison-Wesley Pub. Comp., Reading, MA, 1987.

[Castanon] D. A. Castanon, "Optimal search strategies in dynamic hypothesis testing", IEEE Trans. Sys. Man Cyb. Vol 25, No 7, July 1995 pp 1130-1138.

[Hintz] K. J. Hintz and E. S. McVey, "Multiprocess constrained estimation", IEEE Trans. Sys. Man Cyb., Vol. 21, No. 1, pp 237-244.

[Jenison] Troy Jenison, "Sensor management for multisensor/multitarget detection and classification", Minnesota Center for Industrial Mathematics Master's Thesis, in preparation.

[Kapur] J. N. Kapur and H. K. Kesavan, *Entropy Optimization Principles with Applications*, Academic Press, Inc., Boston, MA, 1992.

[Kastella95] Keith D. Kastella and Wayne W. Schmaedeke, "Method of optimizing the allocation of sensors to targets", Patent Number: 5,451,960; Date of Patent: September 19, 1995.

[Kastella97] Keith Kastella, "Discrimination gain to optimize detection and classification", to appear, IEEE Sys. Man. Cyb., Jan. 1997.

[Mangel] Marc Mangel, "Search Theory: A Differential Equations Approach", in *Search Theory*, David V. Chudnovsky and Gregory V. Chudnovsky, Ed., Marcel Dekker, New York, 1989.

[Musick] Stan Musick and Raj Malhotra, "Chasing the elusive sensor manager", Proceedings 1994 NAECON, Dayton, OH.

[Nash] J. M. Nash, "Optimal allocation of tracking resources", Proc. IEEE Conf. Dec. Cont., 1977, pp. 1177-1180.

[Popoli] Robert Popoli, "The Sensor Management Imperative", in *Multitarget-Multisensor Tracking: Advanced Applications, Vol. II*, Y. Bar-Shalom, Ed., Artech House, 1992.

[Schmaedeke93] W. Schmaedeke, "Information-based sensor management", SPIE Proceedings Vol. 1955, Signal Processing, Sensor Fusion, and Target Recognition II, April, 1993. W. Schmaedeke and K. Kastella, "Information theory and sensor management", Basic Research Group of the Joint Directors of Laboratories 1993 Symposium on Command and Control Research, June 28 - July 1, 1993.

[Stone] Lawrence D. Stone, *Theory of Optimal Search*, 1989, ORSA Books, Arlington, Virginia.

[Wald] Abraham Wald, *Sequential Analysis*, 1947, John Wiley & Sons, Inc., New York, New York.

[Watson] G. A. Watson and W. D. Blair, "IMM algorithm for tracking targets that maneuver through coordinated turns", Proc. SPIE Signal and Data Processing of Small Targets, 1992, Orlando, FL, Apr. 1992, pp 236-247, "Revisit control of a phased-array radar for tracking maneuvering targets when supported by a precision electronic support measures sensor", *Proc. SPIE Signal and Data Processing of Small Targets*, 1994, Orlando, FL, Apr. 1994, pp. 448-459.

[Wetherill] G. Barrie Wetherill, *Sequential Methods in Statistics*, 1966, Methuen and Co Ltd, London, from Methuen's Monographs on Applied Probability and Statistics, and John Wiley and Sons Inc, New York.